APPENDIX

B

ACTIVE DIRECTORY OVERVIEW

Active Directory is a service included with Windows 2000 Server. It stores information about network objects, and makes that information available to Active Directory clients, which can include Windows 2000 Professional as well as Windows 95 or Windows 98 clients. This appendix serves only as an introduction to this topic. To fully understand all the nuances of this service, refer to the Windows 2000 Server documentation and online Help.

In common usage, a **directory** is an information source used to store information about useful, manageable **objects**. For example, a telephone directory stores information about telephone subscribers. In a distributed computing system or a computer network, there are many objects, such as printers, fax servers, applications, databases, and other users. Users want to find and use these objects; administrators want to manage how these objects are used. There is a common need to have usable objects that can be accessed and managed.

In this appendix, the terms directory and directory service refer to the directories found in public and private networks. A **directory service** differs from a directory in that it is defined as both the directory information source (that is, the database of system objects), and the services, such as **LDAP (Lightweight Directory Access Protocol)**, that make the information available to and usable by users and administrators.

What Is Active Directory?

Active Directory is the directory service included with Windows 2000 Server. It extends the features of previous Windows-based directory services and adds entirely new features. Active Directory is secure, distributed, partitioned, and replicated (each of these features is discussed in more detail later in this appendix). It is designed to work well in any size installation, from a single server with a few hundred objects to thousands of servers and millions of objects. Active Directory has many features that make it easy to navigate and manage large amounts of information, generating time savings for both administrators and end users.

Active Directory and the X.500 Standard

Active Directory is based on (but exceeds) the **X.500** standards for directory services. X.500 is a series of International Telecommunications Union (ITU) protocol recommendations that specify a model for connecting local directory services to form one distributed global directory. Following the strict interpretation of the recommendation, local databases hold and maintain a part of the global database, and make directory information available through a local server called a directory service. The user perceives the entire directory to be accessible from the local server.

X.500 also supports data management functions such as the addition, modification, and deletion of entries. Within Active Directory, each directory stores entries in an X.500-like naming scheme, but each directory duplicates the information held by other directories. This duplication provides an implementation that is both faster and more fault-tolerant in a large distributed enterprise.

Active Directory Management

All objects in Active Directory are protected by access control lists (ACLs). ACLs determine who can access the object and what actions each user can perform on the object. The existence of an object is never revealed to a user who is not allowed to access it.

An ACL is a list of access control entries (ACEs) stored with the object it protects. In Windows 2000, an ACL is stored as a binary value called a Security Descriptor. Each ACE contains a security identifier (SID), which identifies the **principal** (user or group) to whom the ACE applies and information on what type of access the ACE grants or denies.

ACLs on directory objects contain ACEs that apply to the object as a whole and ACEs that apply to the individual attributes of the object. This information allows an administrator to control not only which users can access an object, but also what properties those users can see. For example, all users might be granted read access to the e-mail and telephone number attributes for all other users, but security properties of users might be denied to all but members of a special security administrators group. Individual users might be granted write access to personal attributes, such as the telephone and mailing addresses, on their own user objects.

Delegation

Delegation is an important security feature of Active Directory. Delegation allows a higher administrative authority to grant specific administrative rights for **containers** and subtrees to individuals and groups. (See Chapter 5 for more general information on users and groups.) This eliminates the need for domain administrators with sweeping authority over large segments of the user population. For example, there could be a separate administrator for a Corporate Accounting organizational unit.

ACEs can grant specific administrative rights on the objects in a container to a user or group. Rights are granted for specific operations on specific object classes via ACEs in the container's ACL. Specific information on the ACL is covered in Chapter 6. For example, to allow user James Smith to be an administrator of the Corporate Accounting organizational unit, you would add ACEs to the ACL on "Corporate Accounting" as follows:

```
"James Smith";Grant ;Create, Modify, Delete;Object-Class User
"James Smith";Grant ;Create, Modify, Delete;Object-Class Group
"James Smith";Grant ;Write;Object-Class User; Attribute Password
```

Now James Smith can create new users and groups in Corporate Accounting and set the passwords on existing users, but he cannot create any other object classes and he cannot affect users in any other containers (unless, of course, ACEs grant him that access on the other containers).

Inheritance

Inheritance lets a given ACE propagate from the container where it was applied to all children of that container. Inheritance can be combined with delegation to grant administrative rights to a whole subtree of the directory in a single operation. For more information on inheritance, see Chapter 4.

Groups

Windows 2000 introduces the following new group features:

- Groups can be treated as distribution lists (a list of accounts to send e-mail to) if the next major release of Microsoft Exchange (version at this writing is 5.5) is installed.
- Groups can contain nonsecurity members (this is important when the group is used for both security and distribution list purposes).
- Administrators can disable security usage of groups (this is important when the group is solely used as a distribution list).
- Groups can be nested.
- A new group type, the universal group, is introduced.

A universal group is the simplest form of group, i.e., logical components used to associate user objects. Universal groups can appear in ACLs anywhere in the forest, and can contain other universal groups, global groups, and users. Small installations can use universal groups

exclusively and not concern themselves with global and local groups. Additional information about groups is provided in Chapter 5.

A global group can appear in ACLs anywhere in the forest. A global group can contain users and other global groups from its own domain.

A domain local group can be used in ACLs only it its own domain. A domain local group can contain users and global groups from any domain in the forest, universal groups, and other domain local groups in its own domain.

The three group types provide a rich and flexible access control environment, while reducing **replication** traffic to the global catalog (GC) caused by group membership changes. A universal group appears in the GC, but will contain primarily global groups from domains in the forest. Once the global groups are established, the membership in the universal group will change infrequently. Global groups appear in the GC, but their members don't. Membership changes in global groups are not replicated outside of the domain where they are defined. Domain local groups are valid only in the domain where they are defined and do not appear in the GC at all.

Transitive Bidirectional Trust

When a domain is joined to a Windows 2000 domain tree, a **transitive bidirectional trust** relationship is automatically established between the joined-from domain and its parent in the tree. The **trust** is transitive (that is, when a domain joins an existing forest, a trust is automatically established) and bidirectional (that is, the trust exists both ways between the domains, no additional trust relationships are required among tree members). The trust hierarchy is stored as part of the directory metadata in the configuration container (an object that can logically contain configuration information about other objects—for example, a folder is an object that contains file objects).

Active Directory Replication

Active Directory provides **multi-master replication**. Multi-master replication means that all replicas of a given partition are writeable. This allows updates to be applied to any replica of a given partition. The Active Directory replication system propagates the changes from a given replica to all other replicas. Replication is automatic and transparent.

Server Affinity

Windows 2000 uses site information to locate an Active Directory server close to the user to speed communication between the user and the server. A **site** is an area of the network where connectivity among machines is assumed to be very good. Windows 2000 defines a site as one or more IP **subnets**. This is based on the assumption that computers with the same subnet address are connected to the same network segment, typically a LAN or other high-bandwidth environment such as Frame Relay, ATM, or others.

When a user workstation connects to the network, it receives a TCP/IP address from a Dynamic Host Configuration Protocol (DHCP) server, which also identifies the subnet to

which the workstation is attached. Workstations that have statically configured IP addresses also have statically configured subnet information. In either case, the Windows 2000 domain controller (DC) locator will attempt to locate an Active Directory server located on the same subnet as the user, based on the subnet information known to the workstation.

Sites and Replication

The Windows 2000 replication system automatically generates a ring topology for replication among Active Directory servers in a given site. Within a site, directory replication is performed via remote procedure call (RPC). Between sites, replication can be selectively configured to use RPC or messaging. Windows 2000 provides simple Simple Mail Transfer Protocol (SMTP) messaging as a standard feature. If Microsoft Exchange is available, intersite replication can be carried via Exchange, using any of the many mail transports supported by Exchange (this includes SMTP, X.400, and others).

Publishing Objects

Publishing in Active Directory is the act of creating objects in the directory that either directly contain the information you want to make available or provide a reference to the information you want to make available. For example, a published user object might contain useful information about users, such as their telephone numbers and e-mail addresses, whereas a published volume object might contain a reference to a shared file system volume.

When to Publish

Information should be published in Active Directory when it is useful or interesting to a large part of the user community and when it needs to be highly accessible.

Information published in Active Directory has two major characteristics:

- It is relatively static and changes infrequently. Telephone numbers and e-mail addresses are examples of relatively static information suitable for publishing; the user's currently selected e-mail message is an example of highly volatile information.
- It is structured and can be represented as a set of discrete attributes. A user's business address is an example of structured information suitable for Active Directory publishing; an audio clip of a user's voice is an example of unstructured information better suited to sharing via the network.

Operational information used by applications is an excellent candidate for publishing in Active Directory. This includes global configuration information that applies to all instances of a given application. For example, a relational database product could store the default configuration for database servers as an object in Active Directory. New installations of that product would collect the default configuration from the object, simplifying the installation process and enhancing the consistency of installations in an enterprise.

Applications can also publish their connection points (network location) in the directory. Connection points are used for a client/server rendezvous. Active Directory defines an architecture for integrated service administration, using Service Administration Point (SAP)

objects, and provides standard connection points for RPC, WinSock, and COM applications. Applications that do not use the RPC or WinSock interfaces for publishing their connection points can explicitly publish Service Connection Point objects in the directory.

Application data can also be published in the directory, using application-specific objects. Application-specific data should meet the criteria discussed previously. That is, data should be globally interesting, relatively nonvolatile, and structured.

How to Publish

Publishing contents to Active Directory must be handled from Windows 2000 Server. The publishing process is managed by Microsoft Management Console (MMC) snap-ins and other administrative tools that are not included for use with Windows 2000 Professional. From a Windows 2000 server, the means of publishing information varies according to the application or service:

- *RPC:* RPC (remote procedure call) applications use the RpcNs family of APIs to publish their connection points in the directory and to query for the connection points of services that have published theirs.
- *Windows Sockets:* Windows Sockets applications use the Registration and Resolution family of APIs available in WinSock 2.0 to publish their connection points and query for the connection points of services that have published theirs.
- *DCOM:* DCOM (distributed component object model) services publish their connection points via the DCOM Class Store, which resides in Active Directory.

Accessing Active Directory

Active Directory supports clients running Windows 2000 Server and Windows 2000 Professional, as well as Windows 95 and 98 clients that have Active Directory add-on software installed. A client discovers its site by presenting its subnet to the first Active Directory server contacted. The workstation determines the subnet by applying its subnet mask to its IP address. The subnet mask and IP address can be assigned by DHCP or statically configured. The first server contacted uses the presented subnet to locate the Site object for the site where the workstation is located. If the current server is not in that site, the server notifies the workstation of a better server to use.

How Does a Workstation Find a Directory Server?

A workstation finds a directory server by querying the Domain Network System (DNS). Directory servers for a given domain publish Service Resource Records (SRVs) in DNS with names in the following form:

```
_LDAP._TCP.<domain name>
```

Thus, a workstation logging on to *Microsoft.com* might query DNS for SRV records for _LDAP._TCP.killfear.com. A server will be selected from the list and contacted. The contacted server will use the subnet information presented by the workstation to determine the best server, as described in answer to the previous query.

How Do Users Log On?

A user can use a variety of names in a variety of formats to log on to Windows 2000 Professional as part of Active Directory. These include the name formats supported by the Win32 application programming interface **DsCrackNames**, which are used to convert these name forms as necessary.

Domain NetBIOS Name and SAM Account Name

This is the Windows NT 4.0-style logon name. The domain NetBIOS name is the name the domain had prior to migration. The Security Accounts Manager (SAM) account name is the account name the user had prior to migration.

User Principal Name (UPN)

This is in the format <*friendly name*>@<*dotted-dns-domain-name*>. If the name is not unique, the logon attempt will fail with an Unknown User error.

Active Directory and Windows 2000 Professional

As part of an Active Directory, Windows 2000 Professional provides mobile users with the same work environment whether online or offline. As a result, users can work in the same files, folders, or Web sites whether they are connected or disconnected, and easily synchronize those resources. Consistent access to network-based resources helps users stay more productive, whether on an airplane or working at a remote site.

Accessing Files and Folders When Offline

The Offline Files and Folders feature (discussed in Chapter 9) allows mobile users to easily take any combination of files, folders, or entire **mapped drives** with them offline. Instead of using a separate tool, such as the Briefcase, users simply right-click any network-based file or folder and select Make Available Offline from the resulting menu.

When the computer is offline, the files and folders appear in the same directory as they did online—as if they still resided in the same location on the network. This makes them easy to find. Plus, files and folders are visually tagged for offline use by the "roundtrip" arrows in their bottom-left corner.

Synchronization Manager

Using the new Synchronization Manager tool in Windows 2000 Professional, users can synchronize all network resources, including files, folders, e-mail, and databases, in a single location. For details on using Synchronization Manager, see Chapter 15. Users can set the Synchronization Manager to automatically synchronize some or all of their resources. For example, users can set certain files and folders to be synchronized every time they log on or off the network. The Synchronization Manager quickly scans the system for any changes, and if it detects changes, the resources are automatically updated. Only resources that have changed are updated—vastly speeding up the **synchronization** process.

Users can also determine whether files are synchronized when the system is idle, or schedule synchronization for specific time increments, such as every evening. As a result, mobile users always have the most up-to-date information, such as pricing, inventory, or sales data, to communicate to partners and clients—even when traveling.

Synchronization Manager users can also synchronize resources according to their connection types. For example, a user can save time by specifying that large database files only be synchronized when the computer is using a high-speed connection and that all personal documents stored in a specific file are synchronized every time they are connected to the corporate LAN.

Although Synchronization Manager is designed primarily to synchronize documents, it also includes the ability to resolve version conflicts in the event that multiple people edit the same document.

KEY TERMS

Active Directory — A directory and a directory service. The directory is modeled after the X.500 recommendation. In addition, the services component is modeled after the Lightweight Directory Access Protocol (LDAP). Combining these two methods also allows Active Directory to leverage the globally recognized host name resolution protocol DNS (Domain Name System).

container — A logical component used for delegation. Containers contain objects such as "user" type or "computer" type objects.h

delegation — The process of assigning groups or individuals access to manage objects. In Active Directory, delegation allows a domain to be segmented into various logical components. Permissions to manage these logical segments can also be delegated.

directory — An information source used to store information about useful, manageable objects.

directory service — A service that differs from a directory in that it is defined as both the directory information source (that is, the database) and the services (that is, LDAP) that make information available to and usable by the users and administrators.

DSCrackNames — A specific Windows 2000 NTDS API (NT Directory Services Application Programming Interface) that accepts a name and then outputs the desired result. As an example, you could offer DsCrackNames a Windows NT 4 style name "DOMAIN\USER" and request a User Principal Name (UPN). Your result would be *user@domain.com*.

inheritance — A process that lets a given ACE propagate from the container where it was applied to all children of the container.

Lightweight Directory Access Protocol (LDAP) — An X.500-based protocol used to access information directories.

mapped drive — A share on Windows 2000 or NT servers that has been linked to drive letters on the client.

multi-master replication — A replication model that is different from other models because any domain controller can accept and replicate directory changes.

objects — The basis for all things managed in a directory. The "directory dictionary," also known as the schema, defines objects. Objects exist in the form of "user" type, "computer" type. Using these base objects, Active Directory creates an object that can be managed, such as a user "joshuak."

principal — A security object in Kerberos. In Active Directory the Security Principals include Users, Computers, and Groups.

replication (directory replication) — The process of two systems in a homogenous system sharing directory information over the directory services interface. The directory services interface could be based on LDAP or the X.500 DRA (Directory Replication Agent).

sites — The logical definitions in Active Directory that relate to the IP physical substructure of a company. Sites are defined as one or more IP subnets. This in turn relates to your network topology.

subnet — A logical network defined by specifying bits, using the IP addressing and subnetting algorithms (bit anding)

synchronization (directory synchronization) — A process in which two systems in a heterogeneous system share directory information, using an interim agent. The agent contains mapping tables and protocol support for both directories.

transitive bidirectional trust — A standard trust relationship that occurs when a domain joins an existing tree. All domains in the tree have two-way trusts established automatically.

trusts — The administrative links that allow user and group object security information to pass between secure boundaries (domains) in Active Directory.

X.500 — A series of International Telecommunications Union (ITU) protocol recommendations that specify a model for connecting local directory services to form one distributed global directory.